

Porting DRM/KMS drivers to DragonFly BSD

François Tigeot

<ftigeot@wolfpond.org>

About myself

- Consultant, software developer, Sysadmin
- X11 and *BSD user
- Former ccTLD System Engineer
- Introduced FreeBSD in the .fr registry
- DragonFly developer since 2011

About DragonFly


- Unix-like Operating System
- Forked from FreeBSD 4.8 in 2003 by Matthew Dillon
- Aims to be high-performance and scalable
- Uses per-core replicated resources
- Many operations are naturally lockless

About DragonFly (2)

- Innovative features very useful for some workloads
- HAMMER filesystem
 - History retention
 - Deduplication
 - Real-time replication
- Swapcache: second-level file cache
 - Uses existing swap infrastructure
 - Optimized for SSDs

About DragonFly (3)

PostgreSQL 9.3 performance



First tentative

- David Shao, Google Summer Of Code 2010
- Based on Linux + a compatibility layer
- Dropped early but continued in his spare time later
- Apparently got something to work
- Sadly probably specific to his machine
- Only got crashes with my hardware

Starting again with FreeBSD

- Many common kernel APIs with DragonFly, seemed like a good idea at the time
- **June 2012:** Start implementing newly required APIs
- **December 2012:** drm/i915 mostly ported
- **June 2013:** i915 starts working after adding PAT support to the kernel
- **October 2013:** ttm + radeon mostly ported
- **July 2014:** ttm + radeon working reliably

FreeBSD dropped the ball

- No significant updates to drm/i915 after 2012
- Latest supported hardware: Ivy-Bridge
- I switched to Linux as the new i915 upstream in September 2013
- Most important goal: Haswell support

Rebasing on Linux

- Target: Linux 3.8.13
 - Linux 3.8 first version with known working Haswell support
 - drm/radeon port based on Linux 3.8
 - OpenBSD also targeting Linux 3.8.x
 - OpenBSD has the best *BSD drm implementation so far (2014-10)
- Not an easy ride, many issues with the FreeBSD i915 code

FreeBSD issues

- Drm/i915 code vastly different from Linux
- Some very old parts
- Most modern features ~= Linux 3.4
- Two different versions of drm + drm drivers
 - Old drm/ directory
 - New drm2/ directory
- Most code the same, pointlessly duplicated!

FreeBSD issues (2)

- The FreeBSD developer rewrote the entire driver to change its coding style
 - `if (ret) => if (ret != NULL)`
 - `return blah; => return (blah);`
- Different file names
- Functions implemented in different order or present in different files
- June 2014: moved most functions to the same place they were in Linux 3.8.13
- Reduced differences from $\sim 50K$ to $22K$ loc

Update process

- 1. Diff Linux and DragonFly code
- 2. Reduce pointless noise
- 3. Identify important functional differences
- 4. Implement them piecemeal
- 6. Replay Linux git history if required
- 5. Try to keep commits small

Implement Linux APIs

- Same old idea as David Shao's who used a Linux compatibility layer
- Many linux/xxx.h wrappers taken from FreeBSD (OFED project, Infiniband)
- Some APIs implemented locally: `idr(9)`, various `linux/yyy.h` headers, etc...
- Other Linux functions taken from OpenBSD

Implement Linux APIs (2)

- Graphic drivers very complex
- Fast moving targets
- It makes more sense to change the DragonFly kernel to behave like Linux than trying to constantly keep up and change the drivers to use *BSD-specific APIs
- In a way I'm porting DragonFly to the drm drivers and not the drivers to DragonFly

Implement Linux APIs (3)

- `ls /usr/src/sys/drm/include/linux`

<code>atomic.h</code>	<code>export.h</code>	<code>jiffies.h</code>	<code>timer.h</code>
<code>bitops.h</code>	<code>hash.h</code>	<code>kernel.h</code>	<code>types.h</code>
<code>compiler.h</code>	<code>highmem.h</code>	<code>kref.h</code>	<code>wait.h</code>
<code>completion.h</code>	<code>i2c.h</code>	<code>list.h</code>	<code>workqueue.h</code>
<code>delay.h</code>	<code>idr.h</code>	<code>mutex.h</code>	
<code>err.h</code>	<code>io.h</code>	<code>time.h</code>	

- Most common Linux APIs used in i915 ttm and radeon implemented by wrappers
- Reduce differences as much as possible in the `drm/*` directories
- Major exception: locking directives

Difficult spots

- GEM VM code
 - FreeBSD used custom do-everything routines different from the Linux GEM implementation
 - `i915_gem_obj_io()` in FreeBSD
 - Various functions like `i915_gem_shmem_pwrite()` in Linux
- I2c API in Linux, iic API in FreeBSD+DragonFly
- Hard to make these parts of the drivers work like they do in Linux

The road to Haswell support

- 1. Reduce noise
- 2. Replace FreeBSD APIs by Linux ones
- 3. Update the interrupt code
- 4. Update the ringbuffer code
- 5. Update the output management code
- 6. Haswell-specific cache attributes in page tables
- Finally working reliably in August 2014

A few words about Userland

- Pkgsrc packages used until 2013
 - Xorg-server 1.6
 - Required an update
- Switched to Dports
 - FreeBSD ports + adaptation layer + automatic tests and validation
 - Much better
 - Gave us xorg-server 1.12 and contemporary software

A few words about Userland (3)

- Sadly, FreeBSD ports are beginning to lag
- Updating Cairo from 1.10 to 1.12 took years
- Still using xf86-video-intel 2.21.15
- Had to create a locally managed port of xf86-video-intel 2.99.x
- Hope it won't become necessary for other parts

Current state

- drm/i915 mostly in sync with Linux 3.8.13
 - Not the GEM code
- drm/radeon mostly in sync with Linux 3.8
- drm/ttm in sync with Linux 3.9
- Generic drm code a mess
 - A few parts up to Linux 3.8.13
 - Others much much older (< Linux 2.6.26)

Future directions

- Synchronize the drm code itself to Linux 3.8.x
 - Required for key features like DRM Prime, DRM-master, etc...
- Then start upgrading drm+drivers to more recent versions
- Porting drm/nouveau also a good idea
 - Should be easier now that ttm is working properly

Future directions (2)

- Restore 80x25 VGA text mode after Xorg exit
- Currently black screen or frozen image
- Alternatively implement a graphical TTY layer
- Patches are floating around

FreeBSD credits

- Konstantin Belousov
 - Updated drm and drm/i915 to support gem+kms
 - Ported most parts of drm/ttm to FreeBSD
- Alexander Kabaev
 - Started to port the radeon driver
- Jean-Sébastien Pédron
 - Finished the port of drm/radeon from Linux 3.8
 - Critical fixes to drm/ttm

DragonFly Credits

- François Tigeot: bulk of drm, drm/i915, drm/ttm and drm/radeon porting
- Johannes Hofmann: made i915 work
- Matt Dillon: made i915 and ttm work reliably
 - Added PAT support to the kernel
 - fixed critical VM bugs
- Joris Giovanelli and Markus Pfeiffer: critical bug investigations and fixes

Questions ?